PRINCETON UNIVERSITY

CENTER FOR INFORMATION TECHNOLOGY POLICY
PRINCETON UNIVERSITY

December 11, 2019

Before the
Federal Trade Commission
Washington, D.C.

COPPA Rule Review, 16 CFR part 312, Project No. P195404

Thank you for the opportunity to provide comments on the FTC's review of its COPPA Rule to protect children's privacy online. As the online environment evolves and children spend increasing amounts of time on a variety of connected devices, it is important to update the COPPA Rule to keep it current with new privacy risks.

We are academic researchers associated with the Center for Information Technology Policy (CITP) at Princeton University.[1] We write to encourage the Commission to take specific steps to protect and study children's privacy online, drawing on our collective experience in computer science and law. We look forward to further opportunities to engage with the Commission's staff to provide additional analysis as the Rule evolves.

1. **The FTC should revise the COPPA Rule to promote external audits.**

Under the current COPPA Rule, it is rarely possible for expert researchers, civil society groups, journalists, or—most importantly—parents to verify whether a website or online service is respecting children's online privacy. The status quo undermines informed parental decision making and inhibits important forms of oversight.

External auditing has proven invaluable to technology policy in general, and to consumer privacy policy in particular.[2] Studies of online privacy, including studies that

[1] In keeping with Princeton's tradition of service, CITP's Technology Policy Clinic provides nonpartisan research, analysis, and commentary to policy makers, industry participants, journalists, and the public. These comments are a product of that Clinic and reflect the independent views of the undersigned scholars.

[2] There are three primary technical methods computer security and privacy researchers might use to assess whether an online service in is compliance with COPPA: static analysis, dynamic analysis, and network analysis. Static analysis is performed on source code without running it, while dynamic analysis is performed as the code is running. Network analysis is a type of dynamic analysis which involves monitoring incoming and outgoing traffic from an online service.

we have conducted, have repeatedly highlighted questionable practices, prompted corrective changes, and led to enforcement actions. The earliest of these studies emphasized web privacy; more recent work has examined mobile apps, internet of things devices, and most recently, smart TVs. The FTC has consistently elevated this area of research and helped to facilitate a thriving research community, with events like the annual PrivacyCon in Washington and ConPro in San Francisco.

As we discuss further below, the current COPPA Rule inhibits external compliance auditing in two ways. The present "directed to children" approach makes it nearly impossible to know the appropriate compliance tests for any particular website or online service. And the "internal operations" exception poses the same issue, plus makes it nearly impossible to identify whether information flows associated with websites and online services are COPPA compliant. We propose a set of transparency amendments to the COPPA Rule that would address these issues.

A. The FTC should make services disclose whether they consider themselves, in whole or part, "directed to children" under COPPA.

The current COPPA Rule rests on a totality of the circumstances test for whether a website or online service is "directed" to children under 13.[3] If the website or service is not directed to children, the Rule imposes no transparency requirements. If the website or service *would* be directed to children, it may be able to use an age gate to avoid child directed status and associated transparency requirements.

This state of affairs creates foreseeable incentives for websites and online services to reach a determination that they are not child directed, or to implement minimal age gates that are easily gamed. But because the COPPA Rule does not include transparency requirements about a website or online service's child directed status, there is no way to externally distinguish between 1) a website or service that believes it is not covered by the Rule because its content is not child directed; 2) a website or service that believes it would be covered by the Rule but has implemented an age gate such that it is not child directed; and 3) a website or service that believes it is covered by the Rule, but has data

[3] "In determining whether a Web site or online service, or a portion thereof, is directed to children, the Commission will consider its subject matter, visual content, use of animated characters or child-oriented activities and incentives, music or other audio content, age of models, presence of child celebrities or celebrities who appeal to children, language or other characteristics of the Web site or online service, as well as whether advertising promoting or appearing on the Web site or online service is directed to children. The Commission will also consider competent and reliable empirical evidence regarding audience composition, and evidence regarding the intended audience." 16 C.F.R. § 312.2.

practices that fall entirely within the "internal operations" exception (discussed further below).

Each scenario requires a different type of compliance testing: 1) the attributes and audience composition described in the COPPA Rule; 2) the efficacy of the age gate; and 3) the nature of the information flows on the website or service. There is presently no way to definitively know which type of testing is needed and how that testing bears on the provider's overall compliance with the COPPA Rule.

We recommend that the Commission establish a requirement that, if a website or online service determines that it is (in whole or in part) directed to children, it must publicly disclose that fact in a standardized format. In addition to a disclosure for the benefit of human readers, we recommend including a machine readable disclosure to facilitate external auditing. This approach would be technically trivial to implement and would mitigate the current uncertainty about whether websites and online services consider themselves child directed.

We also recommend that the Commission establish a requirement that, if a website or online service is using an age gate as part of its determination that it is not child directed, it must publicly post a description of the operation of the age gate and what steps it took to validate that children under 13 cannot circumvent the age gate. Again, this approach would be straightforward (in both human and machine readable formats), and would consist of merely documenting what websites and online services should already be doing.

As another possible step, we encourage the Commission to consider requiring disclosure from websites and online services that have determined they are not child directed, but acknowledge that they could be plausibly considered child directed. Transparency in these borderline cases would be invaluable for distinguishing websites and online services that are ignoring the COPPA Rule altogether from websites and online services that have at least considered the COPPA Rule and reached a conclusion about its applicability.

B. The FTC should require transparency from providers that rely on the "internal operations" exception.

In the last round of revisions to the COPPA Rule, the Commission established a new exception for collecting and using a "persistent identifier" for the sole purpose of supporting the "internal operations" of a website or online service. While we have

significant reservations about the basis for and scope of the exception—explained further below—we are also concerned about the absence of transparency in how websites and online services are using the exception. As currently implemented, the "internal operations" provision exempts practices from not only verifiable parental consent, but also the Rule's disclosure requirements.[4]

The result of this approach, in combination with permitting the use of first-party and third-party persistent identifiers under the "internal operations" exception, is to prevent meaningful external auditing of how the exception is being used. An audit will reveal that persistent identifiers (typically one or more unique cookies) are being shared with a website or service, and possibly third parties, in connection with activity information (typically a URL). How that activity information is being retained, used, shared, or sold by the website or service—especially to the third parties—will be entirely invisible to external auditing.

We recommend that the Commission establish a requirement that, if a website or service is invoking the "internal operations" exception, it must still publicly disclose its privacy practices, including the specific "internal operations" that each unique identifier will be used for, how long the data will be retained, and which other entities will receive the data.

In revising the transparency dimensions of the "internal operations" exception, we encourage the Commission to consider the "legitimate interests" provision of the European Union's General Data Protection Regulation (GDPR). That provision similarly exempts certain practices from affirmative user consent, but in exchange, an operator must describe the "legitimate interests" that it is relying upon.

2. The FTC should maintain the COPPA Rule's definition of "personal information" and clarify how the definition applies to common practices.

In the most recent COPPA Rule revision, the FTC recognized that "persistent identifiers" are a form of "personal information," because they enable singling out a specific user through their device for contact. This makes sense; we see no basis in computer science for treating persistent identifiers any differently from other means of directing communications, such as telephone numbers or email addresses. While the technical details differ, the use of the information is the same.

[4] "Where an operator collects a persistent identifier and no other personal information and such identifier is used for the sole purpose of providing support for the internal operations of the Web site or online service . . . there also shall be no obligation to provide notice . . . ." 16 C.F.R. § 312.5(c)(7).

We encourage the FTC to make three clarifications to how the COPPA Rule defines "personal information." First, the definition of "persistent identifier" is not quite clear about whether first-party tracking cookies are covered.[5] On the one hand, the definition provides that "a customer number held in a cookie" is covered. On the other hand, the definition requires that a "personal identifier" be usable for tracking a user "across different" websites. Some readers might interpret that qualifier to limit the definition to third-party tracking cookies, even though first-party tracking cookies still enable contact with a user or device, can be reused on other websites (i.e., in a context where they would be third-party tracking cookies), and can be trivially synchronized with third-party tracking cookies. We encourage the Commission to clarify that first-party tracking cookies are included within the definition of "personal information," and to address any special accommodations for first-party tracking cookies in the "internal operations" exception.

Second, we encourage the Commission to clarify that collecting nearby Wi-Fi MAC addresses—a common approach to mobile device geolocation—qualifies as "personal information" under the COPPA Rule. The component of the "personal information" definition that covers street-level geolocation is agnostic to the technology used, but some readers might not understand that network information is routinely translated into geolocation information.

Third, we urge the Commission to clarify that hashed "personal information" is still "personal information." While we recognize that hashing information is a common practice in the online advertising market, the technique does not provide significant privacy protection; it is relatively easy to determine the corresponding personal information that produced the hash.

3. The FTC should revisit the "internal operations" exception to limit its application.

The Commission adopted the "internal operations" exception during the last COPPA Rule revision, in tandem with adding "persistent identifiers" to the "personal information" definition. We understand the purpose behind an internal operations exception since it is a technical reality that, in order to contact a website or online

[5] The discussion accompanying the latest COPPA Rule revisions is also ambiguous, suggesting that first-party persistent identifiers might not be covered as "personal information," but then suggesting that first-party persistent identifiers are covered and addressed by the "internal operations" exception (as we propose). Children's Online Privacy Protection Rule, 78 Fed. Reg. 3,972, 3,980 (Jan. 17, 2013).

service at all, a user must share at least one persistent identifier (their IP address). We also understand that, in at least some circumstances, using a first-party persistent identifier may be necessary to operate a website or online service and may be consistent with child and parent privacy expectations.

We have significant reservations, however, about the scope of the exception. As currently implemented, the exception reflects several distinct purposes for using persistent identifiers: 1) uses that arise from technical necessity (e.g., IP communications require sharing IP addresses); 2) uses that arise from implementing a website or service with generally accepted software engineering and information security practices (e.g., using first-party persistent identifiers to maintain a user session on the website or service or, in some instances, using first-party identifiers to monitor suspicious activity); and 3) uses that are exempt from verifiable parental consent in order to facilitate monetization of websites and online services (e.g., third-party persistent identifiers for contextual advertising and frequency capping).

In reevaluating the "internal operations" exception, we encourage the FTC to clearly articulate the policy underpinnings of the exception. The exception is commonly understood as a provision about technical necessity, rather than as a provision that blends technical, practical, and economic considerations.[6] It is difficult to square that understanding with the current scope of the exception.

We also encourage the FTC to carefully examine whether a parental consent exception for certain third-party advertising practices best effectuates COPPA's legislative framework and is consistent with the Commission's prior positions on online privacy. An important consequence of the current approach is that third-party online advertising firms do not need parental consent to *collect* a detailed record of what children under 13 do online, so long as they do not *use* that record for profiling or behavioral advertisement targeting. The COPPA legislative text establishes a baseline, however, that parental consent is needed for both collection *and* use.[7] And the current

[6] We note that, in the original COPPA Rule, the scope of the "internal operations" provision was much narrower: "those activities necessary to maintain the technical functioning of the website or online service, or to fulfill a request of a child as permitted by § 312.5(c)(2) and (3)." Children's Online Privacy Protection Rule, 64 Fed. Reg. 59,888, 59,912 (Nov. 3, 1999). The exception also had limited effect, merely constraining which information transfers constitute "disclosures" and which entities are "third parties." In essence, the original "internal operations" exception was solely about outsourcing the technical implementation of a website.

[7] 15 U.S.C. § 6502(b)(1)(A)(ii). We also encourage the Commission to evaluate how the current "internal operations" exception squares with COPPA's legislative text. The original version of the "internal operations" exception was essentially an interpretation of the term "disclosure" within the COPPA statute. The current exception, by contrast, goes much further—it creates a new exception to verifiable

COPPA Rule acknowledges that "passive tracking of a child online" is a covered form of collection—exactly the type of third-party conduct that the current "internal operations" provision exempts from parental consent.[8]

Moreover, as the Commission has previously acknowledged, the collection of web tracking and other online activity data has significant privacy implications—it leads to the stockpiling of detailed and sensitive information about consumers—and use-based limitations do not provide adequate consumer control. And yet, at the same time as disclaiming use-based limitations as insufficient for adults, the Commission is relying on a use-based limitation for protecting children under 13—and despite a statutory directive to limit collection.

In addition, to the extent that the advertising components of the "internal operations" exception are grounded in an assumption that certain forms of advertising-related tracking are economically necessary, that assumption might be misplaced. There is limited empirical evidence on the economics of privacy and advertising, and recent research calls into question the comparative bottom line value of advertising-related tracking.[9] We urge the Commission to take this opportunity to facilitate more rigorous study of the economics of privacy and advertising before drawing any firm conclusions about the economic necessity of such practices.

Finally, we urge the Commission to provide greater guidance about when a persistent identifier is "necessary" under the exception, rather than merely a convenient means of implementing specific functionality. In particular, we recommend that the Commission give guidance for the (common) scenario where website or online service functionality *can* be implemented without a persistent identifier or with only a first-party persistent identifier rather than a third-party persistent identifier. We are concerned that, at present, the "internal operations" exception may be viewed as an automatic exception for all first-party and third-party uses of persistent identifiers for covered purposes, and that the qualifier "necessary" is not providing much limitation.

---

parental consent that is distinct from the expressly enumerated exceptions in the statutory text. § 6502(b)(2).

[8] 16 C.F.R. § 312.2.

[9] Marotta et al, *Online Tracking and Publishers' Revenues: An Empirical Analysis*, https://weis2019.econinfosec.org/wp-content/uploads/sites/6/2019/05/WEIS_2019_paper_38.pdf; Davies, *After GDPR, The New York Times cut off ad exchanges in Europe — and kept growing ad revenue*, Digiday, Jan. 16, 2019, https://digiday.com/media/gumgumtest-new-york-times-gdpr-cut-off-ad-exchanges-europe-ad-revenue/.

4. **The FTC should examine the usability of "verifiable parental consent" mechanisms.**

The COPPA Rule relies on verifiable parental consent to allow the collection of personal information of children. But our research into online dark patterns[10] and recent research about consent mechanisms in GDPR[11] indicates that services may not obtain fully informed consent. Specifically, some sites adopt designs that obfuscate the consent mechanism or manipulate users into providing consent that is not entirely freely given or fully informed. It is quite possible that mechanisms for obtaining verifiable parental consent involve similar problematic designs.

Additionally, recent research by our colleagues at CITP indicates that parents' privacy expectations are highly context dependent and contingent on perceptions of the different entities that collect personal information.[12] Depending on the implementation of verifiable parental consent, parents may not fully appreciate which parties will receive data or the consequences of providing approval, and parents might make different decisions were the consent architecture implemented differently.

In light of the growing research literature on privacy choice architecture, we encourage the Commission to examine the usability of common approaches to verifiable parental consent. We recommend that the FTC adopt a data-driven approach to identifying permissible interface designs, and that the Commission offer guidance on best practices when using a permissible design.

5. **The FTC should examine how major platforms can assist with COPPA Rule compliance and verifiable parental consent.**

We encourage the Commission to examine how major platforms, such as Apple iOS/iPadOS and Google Android, could facilitate COPPA compliance. A greater role for the platforms could translate into significantly improved compliance rates and significantly lower compliance costs, especially for smaller developers and content creators.

---

[10] Mathur et al. *Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites.* ACM Conference on Computer-Supported Cooperative Work and Social Computing 2019. https://arxiv.org/pdf/1907.07032.pdf

[11] Utz et al. *(Un)informed Consent: Studying GDPR Consent Notices in the Field.* ACM SIGSAC Conference on Computer and Communications Security *2019.* https://arxiv.org/abs/1909.02638

[12] Apthorpe et al. *Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus COPPA*. USENIX Security Symposium 2019. https://www.usenix.org/system/files/sec19-apthorpe.pdf

One way in which platforms could assist with COPPA compliance is by uniformly flagging users who are under 13.[13] A mobile operating system could, for example, set an under 13 flag when a child account is active on a device. Developers would then have actual notice of the child's age and could take appropriate action (whether implementing COPPA safeguards or automatically limiting access to the app). This approach could mitigate the need for app developers and content creators to implement their own audience management or age gating.

Another way in which platforms could assist with COPPA compliance is by providing verifiable parental consent mechanisms. The major mobile operating systems already provide for linked parent and child accounts; if they also provided a software interface for child accounts to submit permission requests to parent accounts, apps and content could have a convenient and free means of obtaining verifiable parental consent.[14]

These ideas are inherently preliminary, but they are promising, and we encourage the Commission to thoughtfully explore ways to encourage platforms to play a more prominent role in ensuring that apps and content are COPPA compliant.

6. The FTC should study the use of educational technology in the field before considering a specific exception to parental consent.

Providers of educational technology services have made rapid inroads in classrooms and homes across the country in recent years. The Commission asks whether it should consider a specific exception to parental consent for the use of education technology used in schools. We support the Commission's effort to explore this issue, but recommend that further study is needed before developing new rules. Specifically, the FTC should examine how education technology is used in schools (and in homes) and how school administrators select providers and then monitor the information collected, used, or disclosed about students.[15] At minimum, any specific exception to parental consent in the school setting must be accompanied by rules that cover four issues.

---

[13] Another possible direction would be for platforms to implement uniform age gates or audience management features. Our technical intuition is that those directions would be less appealing to platforms and developers, but we mention them here for completeness.

[14] *See* Liccardi et al., *Can apps play by the COPPA Rules?* 2014 Twelfth Annual Conference on Privacy, Security and Trust, https://people.csail.mit.edu/ilaria/papers/LiccardiPST14.pdf; Bélanger et al, *POCKET: A tool for protecting children's privacy online*, Decision Support Systems 54 (2013) 1161–1173, https://www.rsm.nl/fileadmin/Images_NEW/ECFEB/pdf/France_Belanger_supporting_docs.pdf

[15] *See* Alim et al., *Spying on Students: School-Issued Devices and Student Privacy* (EFF 2017), https://www.eff.org/wp/school-issued-devices-and-student-privacy.

First, parents should know what data educational technology providers collect about their children, how that data is used, who has access to the data, and how long it is retained. Whether or not parental consent for such purposes is delegated to a school administrator, the provider should give parents this essential information in an accessible format. Parents should have the right to request that data about their children are deleted, particularly as children transition from grade to grade, or when new educational systems are put in place and old ones are rendered obsolete.

Second, school administrators should be given guidance on how to make informed decisions about selecting educational technology providers, develop policies that preserve student privacy, and train educators to implement those policies. Our preliminary research suggests there is significant work to be done. In one of our studies, the educators in our focus groups understood the need to manage student data responsibly, but school policies sometimes challenged their ability to do so.[16]

Third, the rule should clarify how school administrators and educational technology providers are accountable to parents for how data about their children are collected, used and maintained. In particular, parents should be able to control how educational technology providers use data about children over time to prevent profiling children over multiple years. And, educational technology providers should be prevented from using personal information for marketing purposes to parents and children by default.

Fourth, the scope of educational technology has expanded vastly and can include applications for classroom management (e.g., LanSchool[17]), technology to support home-school connections (e.g. Seesaw[18]), to technology for managing grades (e.g., Schoology[19]), applications to help children learn certain skills (e.g., Dreambox[20]), software to track student behavior and improve disciplinary outcomes (e.g., SchoolFront,[21] BRIM[22]), disability management systems (e.g., AIM[23]), and platforms/devices (e.g., Chromebooks, iPads, Google Classroom[24]). The Commission needs to clearly define what is meant by "educational purposes" in the classroom in considering any exceptions for parental consent.

* * *

[16] Kumar et al., *Privacy and Security Considerations For Digital Technology Use in Elementary Schools*, In CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019). https://doi.org/10.1145/3290605.3300537.
[17] https://lenovosoftware.com/lanschool/
[18] https://web.seesaw.me/
[19] https://www.schoology.com/
[20] https://www.dreambox.com/
[21] https://www.schoolfront.com/discipline-merit
[22] https://antibullyingsoftware.com/
[23] https://www.accessiblelearning.com/
[24] https://edu.google.com/products/classroom/?modal_active=none

As technology becomes increasingly embedded in the lives of children, the FTC's COPPA Rule must evolve to remain consistent with the core statutory purpose of COPPA: to give parents full control over the collection, use or disclosure of personal information about their children. We are available to assist the FTC with analysis that aids this mission.

Respectfully submitted,

Marshini Chetty
Assistant Professor, Department of Computer Science, University of Chicago

Edward W. Felten
Robert E. Kahn Professor of Computer Science and Public Affairs, Princeton University

Mihir Kshirsagar*
Technology Policy Clinic Lead, Center for Information Technology Policy, Princeton University

Arunesh Mathur
Graduate Student, Department of Computer Science, Princeton University

Jonathan Mayer*
Assistant Professor of Computer Science and Public Affairs, Princeton University

Arvind Narayanan
Associate Professor of Computer Science, Princeton University

Victor Ongkowijaya
Graduate Student, Department of Computer Science, Princeton University

Matthew J. Salganik
Professor of Sociology, Princeton University

Madelyn Sanfilippo

Postdoctoral Research Associate, Princeton University

Ari Ezra Waldman
Microsoft Visiting Professor of Information Technology Policy, Princeton University

* denotes principal comment authors.

Contact:

Website: https://citp.princeton.edu
Phone: 609-258-5306
Email: mihir@princeton.edu